Android) and its associated "app store". Using the same platform as friends and family members also eases communication through instant messaging, video conferencing, and photo sharing applications bundled with the operating system.

Many users to stay with the same platform over time, due to brand loyalty, adoption of built-in features like automatic syncing of data with cloud services (e.g., Apple iCloud), and the learning curve for adapting to a new operating system. The users increase the value of their mobile devices through mobile applications, some of which come pre-installed on the device; these applications may also have a significant impact on battery lifetime and bandwidth consumption, though most users have difficulty determining which of their applications are the "resource hogs." Despite the emphasis on the device and the applications, the relationship with the mobile broadband provider is important, too. Most users receive a handset as part of the service contract from their carrier, though the emergence of tablet computers, and changes in the device pricing model being introduced by some carriers (e.g., T-Mobile), are increasing the fraction of mobile devices purchased directly. The mobile broadband provider also has significant influence over the users in terms of pricing plan (e.g., unlimited bandwidth, bandwidth caps, or usage-based billing) and contract restrictions (e.g., early-termination fees, limitations on tethering, etc.).

Application developers: The ecosystem includes a large and diverse group of developers creating applications for a variety of platforms (e.g., Apple iOS and Google Android).

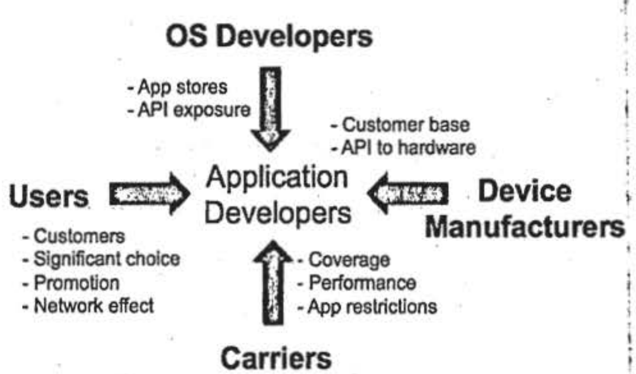

Applications range from network and device utilities, to mobile access to online content, to mobile games, and location-centric applications. Creating a successful application is challenging, and typically requires creating a separate version for each operating-system platform, and relying on whatever Application

Programming Interfaces (APIs) the operating system developers and device manufacturers make available. A range of business models have emerged, as application developers and consumers experiment with different monetization paths, including initial purchase price, "freemium" or free download with limited functionality and pay-to-upgrade charges, ad-supported, and free (or paid) download with in-app purchasing of extras or subscription services. Application developers are somewhat dependent on "App Stores" (the largest app stores are operated by Apple and Google) to distribute their applications, in exchange for a fraction (e.g., 20-30%) of their revenue. In addition, the large number of available "apps" mean that users have tremendous choice, forcing developers to keep prices low to compete with free or low-cost apps; many apps rely on advertising for revenue, and "word of mouth" from users to promote their applications. In addition, application developers rely on mobile broadband providers for good coverage and performance, and are subject to the terms and conditions of the end-user's service contract which may restrict the use of certain apps.

Device manufacturers: Devices such as smart phones, tablets, and smart meters connect to mobile broadband networks. Many end-users identify more strongly with their mobile devices than with their mobile broadband provider.

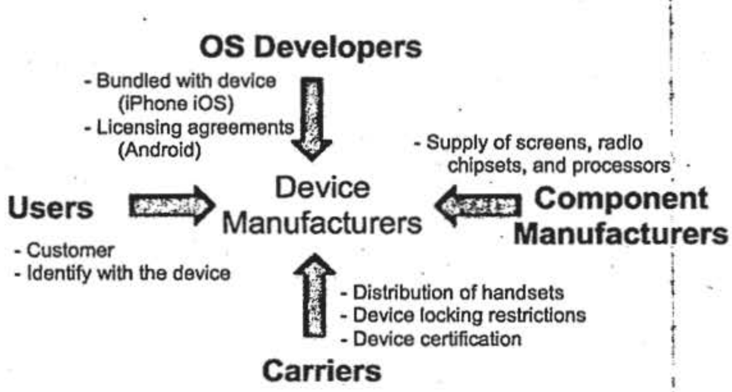

While many handset manufacturers rely on mobile providers to offer sizeable discounts on price of devices sold to consumers (colloquially known as "device subsidies"), the market increasingly includes tablet computers that are sold directly to consumers. Most mobile providers "lock" handsets on their networks, restricting their customers from using the devices with other carriers. The device manufacturers also rely heavily on the component manufacturers for a regular supply of parts. Companies like Qualcomm, Samsung, Intel, and Infineon make radio chipsets and processors that govern radio network operations and compatibility, features, and performance. Even if existing components are limited in functionality, device manufacturers typically find that building their own components is prohibitively expensive. The relationship with component manufacturers is particularly complicated if the company also sells its own mobile devices; for example, Samsung is a leading manufacturer of mobile handsets but is also the primary supplier of screens for its chief device competitor, Apple.

Operating-system developers: The operating system (OS) runs on the devices and provides a development platform for applications. In some cases, the operating system is provided by the device manufacturer (e.g., Apple iOS and Blackberry OS). In other cases, the operating system is provided separately (e.g., Google Android and Microsoft Windows Mobile). Some operating system developers seek to limit the "fragmentation" of the OS software to avoid problems with interoperability, where applications work on one device but not another. Yet, device manufacturers may want to customize the software or experiment with new features. Though Google's Android operating system is open source, recent changes in the terms of service[38] for the Android software development kit prevent developers from creating their own "fork" of the code, to reduce code fragmentation. Similarly, Microsoft's Windows Mobile 7.5/8 is specifically licensed to select hardware partners under terms that greatly limit the variability of the OS implementation across devices. While Android and Apple iOS are by far the largest players in the mobile OS market, the landscape sometimes changes rapidly, as evidenced by the rapid penetration of Google Android OS in the past few years. There are also efforts to launch new, competitive operating systems, such as Mozilla's Firefox OS and Samsung's Tizen. Each OS platform also has very different philosophies towards "openness," with regard to both the OS itself and the application environment it enables.

Mobile broadband providers: Users typically pay mobile carriers to access mobile network services, either through a "post-paid" monthly subscription or a "pre-paid" monthly purchase.

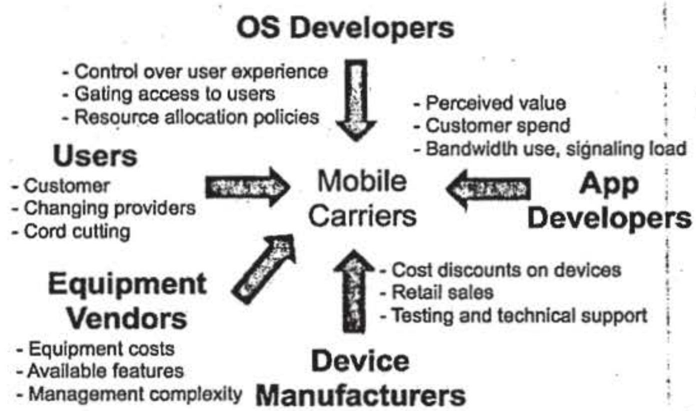

Historically, mobile carriers tightly controlled both the devices and services available to users, but the ecosystem has evolved such that operating system developers, device manufacturers, and application developers have greater control over the user experience and the consumption of network resources. Users who identify primarily with their mobile device may be more willing to change providers at the end of their service contracts, leading to competition over service plans across carriers. The design decisions by application developers influence the consumption of network bandwidth and signaling resource and can degrade performance for all users in congested cells. For example a "chatty" application that sends regular updates every 60 seconds can easily overwhelm signaling resources on the radio access network. The rapid emergence of new applications written by a large community of developers with widely varying expertise makes managing a carrier network challenging. Carriers have little ability to influence a user's choice of applications or an application developer's efficiency in using network resources other than through various forms of usage-based pricing. If data usage continues to grow, carriers will face significant costs to expand network capacity. Carriers' technical options for managing network resources are also limited by the capabilities in the underlying network equipment and mobile devices. Carriers may

[38] http://www.theregister.co.uk/2012/11/15/android_sdk_fragmentation_license_change/

also limit their experimentation with alternative network-management practices to avoid drawing attention from regulators like the FCC.

Network equipment vendors: Mobile broadband providers rely on equipment like cellular base stations, serving and packet gateways, and mobility control software to build and manage their mobile broadband infrastructure.

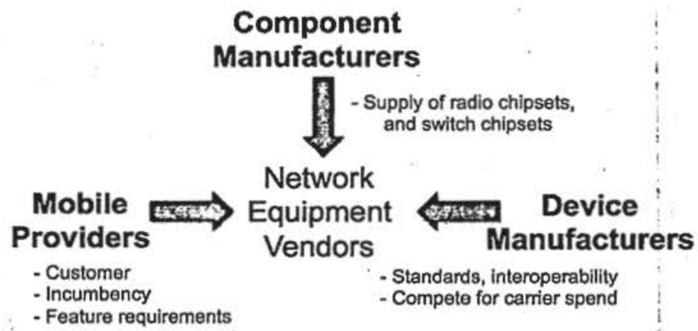

Buying this equipment is a significant capital expense for the carriers as they expand their network footprint, and the capabilities of the equipment influence how the operators can manage their customers' traffic. This, together with the entrance of low-cost players, has driven the rapid commoditization of the network equipment market, and an attendant limit in the level of research and development that can be supported. While the network equipment vendors do not interact directly with end users, or the application and operating system developers, the interplay with device manufacturers is more significant. The network equipment and mobile devices must implement the same standard protocols for the radio access network, leading to cooperation (and competition) in standards bodies leading to more complex standards and the need for extensive interoperability testing. In addition, network equipment vendors must compete with device manufacturers for the limited capital the carriers have to spend on equipment and device subsidies. The network equipment vendors are also dependent on the component manufacturers (e.g., Texas Instruments, Broadcom, and Freescale) that make the chipsets used in their equipment for the radio access and cellular core networks.

In conclusion, the mobile broadband ecosystem has complex power dynamics that affect the incentives each party has to invest in innovation. These dynamics shift rapidly over time in response to business trends (e.g., the prominence of the Blackberry giving way to the iPhone, the emergence of the open Android operating system as a replacement for Apple iOS, and the transition from circuit voice to VoIP with the attendant ecosystem changes). In the next section, we present several case studies of technology and business trends that are affecting openness in the mobile broadband ecosystem.

2. Case Studies

In this section, we present several case studies that illustrate how the inter-relationships in the mobile broadband ecosystem can affect the incentives of different parties to invest and innovate.

2.1 App Stores: App Developers and Operating System Developers

App stores have become an omnipresent feature of mobile broadband. Consumers and app developers both benefit from the convenience that they provide, but app store operators can also restrict the development of mobile applications by leveraging their control over which applications are made available and under what conditions. This section explains some of the motivations for the creation of app stores, explores how app stores may impede openness, and discusses how the trend towards web-based app development might change these dynamics in the future.

The development of mobile app stores – and the app-centric nature of the mobile environment in general – is in some ways a reaction to issues that have arisen with other common software distribution models: traditional desktop software and web-based applications.

In the desktop environment, installed programs have access to a computer's operating system under permission systems that vary as to their robustness and security properties. During the early to mid-2000s, prevalence of malware on personal computers was especially high[39]. The rise in malware was correlated with the emerging prevalence of downloadable, executable content and a runtime model that allowed users to easily and inadvertently introduce malicious code into their machines. Thus, the pure desktop model, with associated malware risks, was seen by some early smartphone innovators as inappropriate for smartphones[40].

Web-based applications, on the other hand, are becoming increasingly robust and are generally safer to run by virtue of the fact that they are confined to the browser[41]. Unfortunately, web applications still lack direct access to many mobile devices' underlying functionality and hardware and thus cannot perform the same functions or provide the same performance as local apps. Although the continued development of HTML5 (discussed below), sophisticated JavaScript APIs, and other web technologies are rapidly pushing web apps forward, in-browser applications still lag behind in some cases in terms of functionality and convenience.

The app-centric model for mobile broadband has therefore been viewed as a way to combine trust and functionality. Apps often undergo review by platform providers and run in a semi-sandboxed environment on the phone's software platform, increasing trust. Because they run locally on the device, they can be hardware-accelerated and have access to a more rich suite of device features than web-based apps.

Apple, Google, Microsoft, and other app store providers shape these dynamics and the overall openness of the mobile app landscape through the policies that they set. These policies concern a variety of technical, operational, and business aspects, including:

- **Installation sources:** On some devices and operating systems (notably Apple's), going through the app store is the *only* way to install an app on non-jailbroken devices. Apple allows web-based applications to be saved as bookmarks, but the user interface and interactions with web

[39] http://download.microsoft.com/download/1/A/7/1A76A73B-6C5B-41CF-9E8C-33F7709B870F/Microsoft_Security_Intelligence_Report_Special_Edition_10_Year_Review.pdf
[40] http://www.nytimes.com/2007/01/11/technology/11cnd-apple.html?_r=0
[41] http://blog.chromium.org/2008/10/new-approach-to-browser-security-google.html

bookmarks and installed apps are not always equivalent. In contrast, Google Play is one of many avenues for app developers to get their apps onto Android devices; Android users can download apps directly from web sites or from other app stores and the OS includes a setting that allows users to "accept apps from unknown sources." Established providers such as Amazon have created their own app stores and developer resources to get apps onto Android-based devices, such as the Kindle Fire.

- **Screening policies:** App store providers have a variety of policies and procedures for screening apps before and after they have been placed in the store. Apps may be reviewed for performance, functionality, access to user data, security, user interface design, and content. Apple reviews all apps before they can appear in the App Store, rejects those that do not meet its App Review Guidelines[42], and may remove apps even after they have been approved. Microsoft uses a similar process and policy[43]. Google generally does not do up-front app screening but removes apps from Play that are found to have security vulnerabilities or that violate Google's terms[44]. Google has also removed specific tethering apps from its app store, reportedly at the request of carriers, because carriers forbid the use of tethering in some of their service plans[45]. Incidentally, the mobile OS vendors also have the capability to remotely uninstall malicious apps[46] directly from users' devices.

- **Revenue-sharing requirements:** App store providers can establish terms that allow them to retain a portion of apps' purchase prices, in-app subscription fees, or ad revenue. Apple, Google, and Microsoft generally retain a 20-30% share of app purchase prices (as does Amazon for its Android-based store)[47]. They may also set the terms about how subscriptions and content can be sold within apps[48].

- **App store navigation**: App store providers choose which apps to feature prominently in their stores and how to categorize apps, at times making decisions that run counter to app developers' desires[49].

All of these policies have the potential to limit the openness of mobile app development. Developers that want to be able to reach users of non-jailbroken Apple devices have no choice but to comply with the terms that Apple sets for the App Store, including the revenue-sharing policies, standards concerning what Apple considers to be "objectionable" content, and technical limitations that include the inability to

---

42 https://developer.apple.com/appstore/guidelines.html
43 http://msdn.microsoft.com/en-us/library/windows/apps/hh694083.aspx
44 http://play.google.com/about/developer-distribution-agreement.html
45 http://news.cnet.com/8301-30686_3-20059461-266.html
46 http://latimesblogs.latimes.com/technology/2011/03/google-removing-virus-infected-android-apps-from-phones-tablets-promises-better-secutiry.html
47 https://developer.apple.com/programs/ios/distribute.html
48 https://developer.apple.com/in-app-purchase/
49 http://www.businesswire.com/news/home/20130314005784/en/Adblock-Reports-Removal-Google-Play-Store-Android

obtain administrative privileges, tether, or alter the "look and feel" of the app[50]. The Android ecosystem is free of many of these limitations, but Google still retains the final say over which apps may appear in Google Play and how easy they are to find and use. On some devices, Google Play is a central source for Android apps despite there being other ways for users to obtain them.

In principle, the convenience and security of the app store model need not be tied to store provider policies limiting the operation or availability of certain apps. Cydia, for example, provides an app store and directory for jailbroken Apple devices, allowing users to more easily discover apps without subjecting app developers to restrictive installation policies or revenue-sharing agreements. While app stores play a pivotal role in the user experience of mobile broadband, it is important to distinguish between the barriers erected by app stores' policies, technical limitations on app development that may be platform-specific but unrelated to app store policies, and the security properties that motivated the development of app stores in the first place. For example, operating system vendors could make the full suite of hardware APIs available to all browsers and apps while still retaining an app store model. This would ease the development of independent apps, but would still subject app developers to the terms set by the app store providers. By the same token, sandboxing and other techniques for making code execution safer could be supported by operating system vendors regardless of whether they enforce an app store model on their platforms or not.

One trend that may shift developers and users away from existing app store models is the continued maturation of the suite of HTML5 technologies[51] [52]. HTML5 comprises the latest versions of the building blocks of the web plus a wide variety of newly developed APIs that give mobile developers access to critical device functionality, including sensors (camera, microphone, etc.), the file system, network interfaces, graphics support, and much more. Because it is based on open, interoperable web standards, the HTML5 technology suite allows developers to build applications from a single code base that work on any device with an up-to-date browser -- which means most any smartphone or tablet already in use. Thus, as HTML5 takes hold as an app development platform, developers will be able to distribute their apps across platforms, independent of whether they are also offered in app stores. HTML5 also includes a variety of security features designed to prevent the kinds of attacks that are often associated with downloadable software and that motivated the development of app stores.

Many HTML5 components are already fully functional and supported by the major browsers, but certain parts of the technology suite are still in the process of being developed and standardized, and questions remain about whether web-based apps can match the performance and user experience of platform-specific ones. As the tools that developers need to create HTML5-based apps that are equivalent or superior to platform-specific apps become increasingly available, the role of app stores in influencing which apps are available and under which conditions may be diminished.

2.2 Service Agreements: Users and Mobile Broadband Providers

Mobile broadband providers have a direct influence on how their customers can access networked services. Service agreements constrain how customers can use their mobile devices. These agreements

---

[50] https://www.eff.org/deeplinks/2012/05/apples-crystal-prison-and-future-open-platforms#gatekeeper

[51] http://www.w3.org/TR/html5/

[52] http://www.whatwg.org/specs/web-apps/current-work/multipage/

illustrate the tensions between the providers' need to limit financial risks (e.g., in discounting or "subsidizing" handsets for customers willing to sign a long-term contract, expanding network capacity, and setting prices for multi-year contracts) and the benefits of giving users flexibility in how they use their mobile devices in a rapidly changing environment.

Billing model: Most mobile broadband providers offer service contracts with a variety of pricing plans. Over the years, unlimited, "all you can eat" data plans have largely given way to plans with bandwidth caps (where subscribers lose network speed after exceeding the cap) or additional charges for additional increments of bandwidth consumption. Still, some providers have many subscribers on "grandfathered" unlimited data plans, increasing the likelihood of high bandwidth consumption when certain applications (e.g., streaming video) or user practices (e.g., tethering) become popular. To manage traffic from these subscribers, some carriers "throttle" top users (i.e., limiting their bandwidth consumption during periods of peak load). Usage caps and usage-based billing encourage users to limit their use of network bandwidth (or defer usage until wired or WiFi connectivity is available), while only indirectly constraining usage during periods of peak load. Alternatives like time-dependent pricing, where providers offer lower prices during off-peak hours (and higher prices when the network is congested), have received significant academic attention, but to our knowledge have not been offered in the market.

Device locking: Many carriers provide customers with a "locked" phone that cannot work with other carriers. Software on the phone ties the subscriber ID (on the SIM card in GSM phones) to the serial number of that particular phone, preventing the customer from using the SIM card in a different phone, or using the phone with a different SIM card. While unlocked phones are common in Europe, most U.S. providers offer locked phones that prevent customers from switching service providers (without buying a new phone), temporarily using a different SIM card during international travel to avoid large roaming fees, or selling an old phone to another user. Providers vary in whether they offer unlocked cell phones (possibly at a higher price) or are willing to unlock a phone after the contract ends (i.e., after recouping the cost of the device subsidy). Recently, the Library of Congress moved to ban mobile users from unlocking their phones without the carriers' permission[53], treating attempts to circumvent device locking as violating the anti-circumvention provisions of the Digital Millennium Copyright Act (DMCA). In response, some regional and rural providers have supported efforts to allow users to legally unlock their phones[54] without their providers' permission.

Tethering: Many providers restrict customers from "tethering" to share a mobile broadband connection with other devices, such as a laptop. Some providers do not allow tethering on certain data plans (e.g., unlimited plans), or require customers to pay extra (above the normal cost of their data plan) for tethering. The rationale is that tethering often leads to a substantial increase in bandwidth usage, beyond what the provider may have anticipated when designing its network and pricing structures. In 2012, Verizon was accused of requesting that Google remove tethering applications from the Android app store, so customers could not use these applications as a way to avoid paying a $20/month tethering fee. The FCC ultimately reached a consent decree[55] and settlement with Verizon, under the terms of which Verizon

[53] https://dl.dropboxusercontent.com/u/3155588/SJUD%20cell%20phone%20bill.pdf

[54] http://www.mobilenapps.com/articles/7901/20130314/phone-unlocking-small-carriers-backing-bill-for-apples-iphone-access.htm

[55] http://www.fcc.gov/document/order-and-consent-decree-verizon-wireless-pay-125-million

could not block access to tethering applications[56], making it possible for users with unlimited data plans to tether without paying extra charges; customers subject to usage caps or usage-based billing would have their tethering traffic metered just like any other data traffic. This decision by the FCC was specific to Verizon (under the conditions attached to spectrum licenses that Verizon purchased at auction), and the FCC has not taken any action as to other providers.

Application restrictions: Some providers impose restrictions on what mobile applications a subscriber can run under specific pricing plans. A good example is the evolution of AT&T's policies concerning Apple's FaceTime application for high-quality video calls, as discussed in an earlier report[57] from our OIAC working group. FaceTime is automatically integrated into the calling features of the mobile device, and makes heavy use of radio network bandwidth in both directions between the device and the cellular base station. When FaceTime first became available over cellular data networks, AT&T limited the use of FaceTime to customers of its MobileShare data plan, where multiple devices share a single limit for total data usage. Later, AT&T broadened the range of plans that support FaceTime, but still did not support the application for subscribers on its legacy unlimited data plan; recently, AT&T announced that all customers[58] (even those on unlimited data plans) will be able to run FaceTime over the cellular LTE network by the end of 2013. Another example of carriers imposing application restrictions occurs when they prohibit the use of tethering applications in their terms of service. These restrictions sometimes arise *after* a customer has chosen a specific service contract, when the emergence of a new application leads to heightened concerns about sudden increases in bandwidth usage.

Two-sided pricing: Usage caps and usage-based billing naturally make users conservative about running bandwidth-intensive applications (e.g., video streaming and online gaming). Some content providers and mobile providers may be willing to offer "toll free" or "sender pays" services, where the bandwidth consumed is sponsored or paid by the content provider, rather than counted towards the customers' usage cap. Broad use of two-sided pricing is not (yet) common in the U.S. mobile broadband market[59], though several European and Asian providers have partnered with content providers to offer plans that do not count applications like Facebook and Spotify against a usage cap[60]. These trends raise interesting questions about openness. On the one hand, "toll-free" data may facilitate end-users' ability to access mobile content at a reasonable cost from those providers willing to subsidize the cost of delivering the data. Enabling content providers to pay for data delivery offers users an incentive to access the sponsored content. In the short run, this is beneficial for consumers of that content, particularly for budget conscious users on smaller data plans. On the other hand, sponsored delivery potentially works *against*[61] the goals of

---

[56] http://bits.blogs.nytimes.com/2012/07/31/fcc-verizon-tethering/

[57] http://transition.fcc.gov/cgb/events/ATT-FaceTimeReport.pdf

[58] http://www.macobserver.com/tmo/article/att-opening-facetime-over-cellular-to-all

[59] Discussions of two-sided pricing sometimes reference the Amazon Kindle e-reader device, which in some cases is sold to users without requiring them to purchase a separate service contract with a carrier despite the fact that the device uses a cellular network. However, e-book downloads consume relatively little bandwidth and do not constitute general, universal Internet service. As the Kindle started supporting basic Web browsing, and some users started tethering the device to use as a mobile hotspot, Amazon started capping the free cellular bandwidth usage to 50 megabytes per month.

[60] http://www.npt.no/marked/markedsregulering-smp/marked/marked-7/_attachment/2362?_ts=139b9fde471

[61] http://media.law.stanford.edu/publications/archive/pdf/schewick-statement-20100428.pdf

openness because (i) increasing the costs for content providers may reduce innovation and (ii) smaller, upstart content providers cannot easily amortize the "chargeback" costs through advertising revenue or subscription fees. Entrenching the largest content providers that have the means to strike deals for sponsored data with carriers puts new entrants at a disadvantage. This is clearly an area of ongoing debate.

The evolution of service contracts and pricing plans show that there is a great deal of experimentation in mobile business models, which is enabling innovation and value to customers and others in the ecosystem. Some business models raise concerns about carriers restricting the way consumers use their mobile devices and about long-term impacts on application and content innovation.

2.3 Network-Unfriendly Apps: Mobile Broadband Providers and App Developers

The applications running on mobile devices have a profound influence on the network resource demands for mobile providers. While supporting the resource demands of applications is also important in wireline networks, mobile broadband networks raise several unique challenges. First, mobile apps are written by millions of software developers, including an unprecedented number of novice programmers who have little understanding of how high-level design decisions affect the usage of network and battery resources. Second, radio access networks have very limited bandwidth, particularly on the "uplink" from the mobile devices to the cellular base station, making it relatively easy for one rogue application to consume most of the available resources. Third, communication in cellular networks requires mobile devices to first establish a "bearer" with the base station, leading to signaling overhead. Fourth, expanding the capacity of a cellular network requires a substantial upfront investment for acquiring spectrum licenses, deploying cell towers, and transitioning to new technologies (e.g., LTE).

For mobile providers, applications that (unwittingly) consume excessive bandwidth and signaling resources cause congestion for other users in the short term, and require a larger investment in network capacity in the long term. In addition, applications that waste network bandwidth or battery lifetime limit the value of a mobile broadband service to end users, particularly if users are subject to usage caps or usage-based billing. As a result, without greater transparency to increase user awareness of an application's efficiency -- and usage-based pricing models to incent them to choose the most efficient applications -- providers could see a limited return on the substantial investment required to expand network capacity, and still face the risk of a new mobile application swamping the available resources. Mobile applications can consume excessive network resources in several ways:

Chatty applications consuming excessive signaling resources: In contrast with wireline networks, mobile devices cannot communicate over a cellular network without first establishing a "bearer" to the cell tower. Establishing a bearer requires the mobile device to exchange several control messages over the cellular network. To avoid the overhead of establishing a new bearer, the mobile device continues to occupy transmission channels and codes until a period of inactivity expires. As such, transmitting a small amount of data can consume significant resources in the radio access network, as well as significant battery resources on the mobile device. The problem is exacerbated by "chatty" applications that periodically send short messages to monitor user behavior, maintain a connection for "pushing" data to the mobile device, or update the display of advertisements. Depending on the frequency of the messages, each transmission may require establishing a new bearer, at the expense of additional signaling resources.

A recent study[62] showed that some applications consume as little as 1.7% of network bandwidth, but up to *30%* of signaling capacity. Signaling load is a low-level issue that even a seasoned application developer might not consider, and it may cause an application that worked perfectly well on a wireline network to overwhelm a cellular network.

Aggressive applications consuming excessive bandwidth: The Internet relies on end-host computers to adapt their sending rates in response to network congestion, to ensure fair sharing of the available bandwidth. Applications using the Transmission Control Protocol (TCP) automatically send data more quickly when the network is lightly loaded, and more slowly when the network is congested enough to drop packets. In addition to decreasing the sending rate, multimedia applications may adjust the audio or video encoding to continue streaming data quickly enough for continuous playback despite the reduced available bandwidth. However, some applications do not use TCP or perform "TCP-friendly" congestion control, open multiple parallel TCP connections to receive a larger share of the limited bandwidth, or do not use adaptive content encodings. The encoding issue was apparently at play with Apple's FaceTime application, as discussed in an earlier report[63] by this OIAC working group. In addition, some operating systems are intentionally more aggressive than the protocol standards prescribe in sending data at the start of a TCP connection[64], to reduce latency particularly for small transfers. Given the Internet protocols place important resource-management functionality at the end hosts, the sharing of the limited bandwidth in a cellular network is not completely within the provider's control.

Inefficient applications transferring redundant data: A mobile application often needs to display the same data to the end user more than once, such as previously-downloaded images or articles. Caching content on the mobile device is an effective way to avoid duplicate transmission of the same data, reducing the consumption of battery, bandwidth, and signaling resources. Despite some support for caching on mobile devices, a recent study[65] found that redundant data transfers still consume 18-20% of bandwidth and 6% of signaling load. Rather than performing data transfers themselves, many mobile applications use HTTP libraries. Unfortunately, many of these libraries do not perform caching at all, or do not fully support the HTTP protocol standards for caching. Similarly, some mobile Web browsers do not make effective use of caching. In addition, cached data does not always survive an application crashing or a mobile device rebooting, leading to further wasted transfers and battery resources. In some cases, software bugs can cause excessive downloading of redundant content, as was in the case with an earlier bug in Apple iOS 6.0[66] that caused duplicate downloads of certain podcasts[67]. Enforcing usage caps and usage-based billing can help carriers recoup the cost of duplicate data transmissions, but also gives users the perception of a lower quality of experience for a given price for their mobile broadband service, and does not provide a direct incentive to app developers to reduce redundant transmissions.

---

[62] Feng Qian et al, "Periodic transfers in mobile applications: Network-wide origin, impact, and optimization," in *Proceedings of the World Wide Web Conference*, May 2012. http://web.eecs.umich.edu/~zmao/Papers/periodic_www2012.pdf

[63] http://transition.fcc.gov/cgb/events/ATT-FaceTimeReport.pdf

[64] http://blog.benstrong.com/2010/11/google-and-microsoft-cheat-on-slow.html

[65] Feng Qian et al, "Web caching on smart phones: Ideal vs. reality," in *Proceedings of MobiSys*, June 2012. http://web.eecs.umich.edu/~zmao/Papers/caching_mobisys2012.pdf

[66] http://venturebeat.com/2012/11/14/ios-6-0-bug-causing-massive-data-consumption-on-podcasts/#bmb=1%20%E2%80%A6

[67] http://labs.prx.org/2012/11/14/ios-6-0-devours-data-plans-causes-cdn-overages/

Although applications may consume excessive resources, the incentives of all of the parties---application developers, mobile broadband providers, and end users---are generally aligned. More efficient applications lead to better performance (and better battery lifetime) for users, and lower loads on the network. As such, the main challenges are *education* (of application developers, so they can write network-friendly apps) and *visibility* (for users, so they know which applications are hogging resources). A good example of education of developers is AT&T's Application Resource Optimization (ARO) tool[68] and associated training, which helps application developers understand how their apps would behave on mobile broadband networks. ARO helped the developers of the popular Pandora application substantially reduce their consumption of energy and signaling resources by transmitting audience measurement data less frequently. A good example of visibility is the reviews of applications in app stores, which increasingly comment on an application's use of battery and bandwidth (though not signaling load). Further investment in tools, training, and rating of applications would help application developers and users alike make more informed decisions about resource consumption.

2.4 Wi-Fi Offloading: Competition for Mobile Providers

One technology trend that is changing the dynamics of the mobile broadband market is the growth of non-commercial, wireless Internet access, typically provided using unlicensed spectrum approaches such as Wi-Fi, in many cases, backhauled over a pre-existing (wired) broadband connection.

Over the past 10 years, there has been an exponential growth in cellular data traffic, driven primarily by the dramatic increase in use of smart phones and tablets. As a consequence of the growth in demand, mobile broadband providers are aggressively expanding their network capacity. In addition, due to the prevalence of Wi-Fi on smart phones and tablets, and the increasing availability of Wi-Fi-enabled Internet service in public places (e.g. coffee shops, airports, campuses, hotels) and Wi-Fi-enabled routers at home and in the enterprise, an increasing fraction of mobile wireless data traffic is carried over Wi-Fi access, rather than cellular networks, with different studies suggesting that anywhere from 20-80% of wireless data traffic is carried over Wi-Fi, and ~30-50% of the 'mobile' data traffic may be *cost-effectively* offloaded from cellular networks, depending on the specific deployment scenario[69].

One of the key differences between Wi-Fi networks and cellular networks is that Wi-Fi users may be subject to interference from users of neighboring access points. The quality of a Wi-Fi connection as compared to a cellular data connection may therefore suffer in the presence of interference due to a lower signal-to-noise ratio, resulting in a significantly diminished throughput relative to cellular networks in public settings; a recent paper[70] suggests that less than a third of mobile data traffic may be carried over Wi-Fi networks even in campus environments with dense Wi-Fi deployments. Likewise, similar Quality of Service (QoS) mechanisms that offer hierarchical or differential scheduling and queuing of data flows

---

[68] http://www.att.com/gen/press-room?pid=22388

[69] Randall Schwartz and Magnus Johansson, "Carrier WiFi Offload: Building a business case for carrier WiFi offload," Wireless 20/20, March 2012. http://www.wireless2020.com/docs/CarrierWiFiOffloadWhitePaper03202012.pdf

[70] Shu Liu and Aaron Striegel, "Casting doubts on the viability of WiFi offloading," in *Proceedings of ACM SIGCOMM Workshop on Cellular Networks*, August 2012. http://conferences.sigcomm.org/sigcomm/2012/paper/cellnet/p25.pdf

with different priorities may not be available on Wi-Fi and cellular connections, depending on their configuration. But the availability of cheap or free capacity (and considerable spectrum, e.g., ~400 Mhz in the 5Ghz band[71]) makes Wi-Fi-based solutions attractive for simple web services delivery. Furthermore, with the emergence of usage-based pricing for cellular data services, which encourages users to manage their cellular data usage, and provides unlimited access when the user is connected to certain Wi-Fi Access Points (their own at home, or in a public place), it is legitimate to ask "will Wi-Fi eventually carry a large enough share of mobile user traffic to cause a significant change in the mobile broadband market, and change the essential economics?". This section explores some aspects of this question.

To address this question, we must first identify the types of Wi-Fi solutions. For the purposes of this discussion, we characterize three types of Wi-Fi: (i) non-public indoor (owned/operated by an individual or business), (ii) public indoor including both free or fee-based (likely owned and operated by a business, and provided to its customers) and commercial (owned and operated by a Wi-Fi network operator), and (iii) public outdoor (likely owned and operated by a network provider or campus-based business, or municipality).

These different types of Wi-Fi access points have different characteristics in terms of accessibility, security, and performance, as well as different degrees of utility to the user. They also have different economics. The benefits and limitations of each are summarized in the following table:

	Cost to Operate	Accessibility	Service Continuity	Radio Performance	Commercial Service	Cellular Offload Potential
Type 1 (non-public indoor)	**Low** (unmanaged & connected to existing BB)	**Limited** (only to individual users or employees)	**Limited**	**Not managed**	**No**	**> 50%**
Type 2 (public indoor)	**Medium** (managed by connected to existing BB)	**Good** (subject to business rules)	**Some** (indoor continuity)	**Some management**	**Yes** (direct or indirect payment or subscription)	**< 50%**
Type 3 (public outdoor)	**High** (managed & uses new network connection)	**Good** (subject to subscription or business rules)	**More** (outdoor continuity and cellular networking)	**More management**	**Yes** (subscription service)	**< 50%**

The preceding table summarizes the essential properties of the different Wifi deployment types, with the table categories and entries defined as follows:

- Cost to Build and Operate: This refers to existence of a backhaul network, power, and Wi-Fi access point management

[71] http://en.wikipedia.org/wiki/List_of_WLAN_channels

  - Low cost: Pre-existing, economical backhaul and power with no AP management
  - High cost: New backhaul and power network and sophisticated management

- Accessibility: This refers to the ability to connect to Wi-Fi APs
  - Limited: Restricted only to certain users (e.g. employees)
  - Good: Can be accessed by anyone willing to subscribe or agree to terms and conditions

- Service Continuity: This refers to ability to maintain a session or connectivity when moving from one location to another
  - Limited: Little or no ability to seamlessly connect to neighboring AP
  - Some: Able to maintain session between APs in similar location, from same provider
  - More: Session and service continuity by interworking with other APs and/or the cellular network

- Radio Performance: This refers to management of the Wi-Fi air interface
  - Not managed: Air interface configuration independent of all other APs
  - Some management: Some coordination of APs via common controller
  - More management: Coordination of APs via common controller, with interference management

- Commercial Service: This refers to whether a Service Provider owns and manages APs
  - No: APs owned by private individual or entity
  - Yes: APs owned by commercial entity (business, building provider) or Service Provider

- Cellular Offload Potential: This concerns the potential of a Wi-Fi AP to offload cellular network traffic
  - There are many different estimates of the how much data offload can be achieved by a Wi-Fi network (see the preceding references for examples), but it is broadly agreed that somewhere between 50-75% of time the average user is in home or in an enterprise environment where Wi-Fi experiences relatively little interference and so is highly effective at offloading data traffic, and consequently only 25-50% of the time is the user outdoors or in a public indoor location, where a combination of Wi-Fi and cellular networks would provide the solution.

What does this simple analysis suggest about the impact of Wi-Fi solutions on the mobile broadband market? The growth of these free or lower cost alternatives in any market clearly benefits consumers in terms of providing access to more wireless capacity. However, it is also the case that the user experience amongst Wi-Fi services varies widely, with registration procedures not being seamless, the network performance sometimes poor due to interference, and inconsistent deployment of recent Wi-Fi security enhancements. Some of these issues are being addressed by the Hotspot 2.0 initiative[72] of the Wi-Fi Alliance, which seeks to increase the degree of 'management' of Wi-Fi access points, and to provide seamless authentication and session continuity (between Wi-Fi access points within the same area). Based on these trends, mobile operators are increasingly integrating Wi-Fi solutions with their cellular offers and encouraging use of Wi-Fi for unlimited data offload for 'best effort' services. Indeed, 3GPP is working in standards to allow seamless session continuity between cellular and Wi-Fi solutions, per serving area or per cell, or even per application in future, based on the local availability of capacity, and the needs of the application, as well as user preference and services agreements.

[72] http://www.cisco.com/en/US/solutions/collateral/ns341/ns524/ns673/white_paper_c11-649337.html

These emerging trends effectively mean that Wi-Fi will not just be a wireless broadband solution, but will also become an essential part of providing mobile broadband services to users. Furthermore, given the lower barrier for entry into the Wi-Fi solution space (due to the absence of the need to acquire spectrum or to support wide-area coverage, or mobility), the number of providers that can and will likely enter this space is significant and will likely therefore stimulate additional innovation in wireless data services.

So the future of mobile broadband should consider the combined roles of licensed and unlicensed spectrum solutions, as they are complementary parts of the space, with licensed spectrum approaches providing coverage and capacity with full mobility, security, and quality of service, and unlicensed approaches providing additional capacity with some (e.g., indoor) mobility and nomadicity, but with more limited QoS capabilities and inconsistent security implementation, at least in the near future.

Looking forward, there will be further evolutions of this licensed/unlicensed paradigm to include 'shared spectrum' approaches, based on white-space spectrum (spectrum in and around the TV frequencies that is either unused or infrequently used) or in higher frequency bands such as the 3.5GHz band currently licensed for military use, but for which the FCC has indicated the desire to make available for commercial use by multiple parties in a shared way (use it when you need it, then release it) in a Notice of Public Rule Making (NPRM)[73].

Consequently, we conclude that the user mobile broadband experience will be provided by a combination of complementary approaches, and potentially a variety of different providers, indoor, outdoor, at home, and at work. This dynamism to the mobile broadband market suggests that the future of user choice and experience delivery will continue to grow and expand, with increasing value delivered by the expanded ecosystem.

3. Conclusions

The mobile broadband ecosystem is complex and dynamic, with a variety of players affecting the user experience and the incentives for further innovation and investment. This report encourages the FCC to take a broad view of interactions between the different players in the mobile broadband ecosystem, even though most of the parties involved are not subject to the Open Internet Order. Also, we recommend being watchful of recent trends, such as HTML5 and Wi-Fi offloading, that may lead to greater competition, as well as the emergence of several "vertical players" with growing influence spanning multiple parts of the ecosystem.

We believe that transparency, education, and competition are important complements to existing FCC oversight in helping achieve the goal of a healthy mobile broadband ecosystem. Transparency can take many forms, such as the disclosures required by the Open Internet Order, and improved communication to users (about applications' battery and network resource consumption) and application developers (about the policies by which app stores and carriers might restrict access to their applications). Education includes teaching application developers how to create "network friendly" applications. Finally, competition includes both a healthy balance *between* the various parts of the ecosystem as well as having

[73] http://www.fcc.gov/document/enabling-innovative-small-cell-use-35-ghz-band-nprm-order

multiple viable choices *within* each part of the ecosystem. The combination of all these factors will help ensure all players – not just those subject to the Open Internet Order – contribute to the openness and health of the mobile Internet.

FCC Open Internet Advisory Committee Mobile Broadband Working Group

- Chair: Jennifer Rexford, Professor of Computer Science, Princeton University
- Harvey Anderson, Vice President of Business Affairs & General Counsel, Mozilla
- Brad Burnham, Founding Partner, Union Square Ventures
- Alissa Cooper, Chief Computer Scientist, Center for Democracy & Technology
- Charles Kalmanek, Vice President of Research, AT&T
- Matthew Larsen, CEO, Vistabeam
- Dennis Roberson, Vice Provost & Research Professor, Illinois Institute of Technology (representing T-Mobile)
- Chip Sharp, Director, Technology Policy and Internet Governance, Cisco Systems
- Marcus Weldon, Chief Technology Officer, Alcatel-Lucent
- Jonathan Zittrain, Professor of Law and Professor of Computer Science, Harvard; OIAC Chair (ex officio)
- David Clark, Senior Research Scientist, MIT Computer Science and Artificial Intelligence Research Laboratory; OIAC Vice Chair (ex officio)

Specialized Services: Summary of Findings and Conclusions

FCC Open Internet Advisory Committee
Summary of findings and conclusions, July 2013

The Specialized Services working group prepared a series of case studies to explore issues in the specialized services landscape, and created a series of conclusions based on those case studies.

The Open Internet Report and Order (R&O) assigned to the Open Internet Advisory Committee ("OIAC") the task of aiding the FCC in the task of monitoring specialized services for their impact on Internet access.[74] As part of the proceedings of the Open Internet Advisory Committee, the Specialized Services working group has met for the 12 months prior to the July 2013 meeting of the committee. This report summarizes the findings and conclusions of the working group.

We organized our work around two tasks:

- Attempting to articulate a careful definition of the term "specialized services", and considering whether the working group has advice to the FCC on the criteria that will prove useful in practice to define and characterize a specialized service.
- Developing advice to the FCC with respect to how they should monitor the impact of specialized services on the character of broadband Internet access service (BIAS).

Background

The ability to offer multiple services was an initial driver for many of the significant network investments made by service providers in higher capacity broadband access network architectures. For legacy telephone operators, the emergence of VDSL and ADSL2+ and MPEG-4 enabled them to leverage their existing copper infrastructure to more rapidly deliver a "triple play" of services: voice, data, and video. Similarly, the cable operators have used their platform to deliver a range of services. The current trend is that all these services will migrate to a provider platform based on the Internet protocol (IP). The R&O uses the term "specialized services" to identify those IP-based services that are not subject to the FCC's Open Internet rules.

The R&O states that the specialized services category in the report could raise two concerns that it would monitor going forward. First, the FCC should guard against the possibility that a broadband provider might label a service as a specialized service that would otherwise be correctly identified as an Internet access service in order to evade Open Internet rules. Second, broadband providers might constrict or fail to continue expanding network capacity allocated to broadband Internet access service in order to provide relatively more capacity for specialized services.

[74] Preserving the Open Internet Broadband Industry Practices, GN Docket No. 09-191, WC Docket No. 07-52, FCC 10-201, 114 (Dec. 23, 2010) [hereinafter *R&O*].

The FCC notes that their goal is to achieve a balance of innovation in infrastructure and applications, but the report does not state any conclusions as to the impact of specialized services on that objective. On the one hand, the R&O notes that: "specialized services may raise concerns regarding bypassing open Internet protections, supplanting the open Internet, and enabling anticompetitive conduct."[75] The advantages to a facilities owner of deploying a service as a specialized service, as opposed to an OTT service, is that the facility owner can offer the service with attributes such as a guaranteed quality of service not permitted today with BIAS, and thus not accessible to competitive OTT services[76]. On the other hand, the benefits to the consumer of specialized services are considerable. The business case to justify the investment in the expansion of fiber optics and improved DSL and cable technology which led to higher broadband speeds was fundamentally predicated upon the assumption that the operator would offer multiple services: while all offerings present uncertainty and risk, the projected value that consumers placed on multiple offerings promised an acceptable return on the investment in the expansion of the overall broadband infrastructure, while the value consumers placed on increased BIAS speeds alone did not yield acceptable projected returns.[77] This appears to remain true today, as even new entrants such as Google Fiber offer video services in addition to BIAS[78]. Accordingly, high speed internet access service has benefited from the deployment of specialized video services like IPTV, because the investment in the higher bandwidth infrastructure needed for video services brought higher capacity to more households.

Defining specialized services

Our starting point in this discussion was to see if we could agree on a meaning of the term "specialized services", as given to us by the FCC. This proved difficult. The Open Internet Report and Order defines a specialized service as a service "that broadband providers may offer... over the same last-mile connections used to provide broadband service."[79] Examples of specialized services mentioned in the R&O include facilities-based VoIP, IP video,[80] e-reading services, heart rate monitoring, and energy sensing.[81]

The use of the term in the R&O is in the context of the scope of the rule-making, which is set forth as following[82]:

"We find that open Internet rules should apply to "broadband Internet access service," which we define as:

---

[75] *Id.* at 112.

[76] Independent of whether it is in the business interest of a BIAS provider to offer QoS, the R&O may not permit this option.

[77] The FCC has concurred with this assessment in its Report and Order relating to local cable franchising: see In the Matter of Implementation of Section 621(a)(1) of the Cable Communications Policy Act of 1984 as amended by the Cable Television Consumer Protection and Competition Act of 1992, MB Docket No. 05-311, FCC 06-180, para 51.

[78] For a discussion of the role of video in the Google fiber offering, see http://news.cnet.com/8301-1023_3-57586894-93/google-exec-sees-google-fiber-as-a-moneymaker/

[79] *Id.* at 7.

[80] *Id.* at 61.

[81] *Id.* at 33.

[82] *Id.* at 44.

> *A mass-market retail service by wire or radio that provides the capability to transmit data to and receive data from all or substantially all Internet endpoints, including any capabilities that are incidental to and enable the operation of the communications service, but excluding dial-up Internet access service. This term also encompasses any service that the Commission finds to be providing a functional equivalent of the service described in the previous sentence, or that is used to evade the protections set forth in this Part."*

With some informal guidance from the FCC, the working group took as a starting point that the term "specialized services" describes anything not covered by this rule. In other words, the group took the term to describe services that are "anything else". This inclusive definition would imply that for purposes of the R&O, the category of specialized services would include services regulated in other ways by the FCC, including voice and video.

However, this inclusive definition proved very difficult for the working group to accept in our discussions, because the term has also been used by the FCC elsewhere in less inclusive ways. The R&O itself refers to specific text in the Open Internet NPRM, which defines specialized services as follows:

"As rapid innovation in Internet-related services continues, we recognize that there are and will continue to be Internet-Protocol-based offerings (including voice and subscription video services, and certain business services provided to enterprise customers), often provided over the same networks used for broadband Internet access service, *that have not been classified by the Commission*. We use the term "managed" or "specialized" services to describe these types of offerings. The existence of these services may provide consumer benefits, including greater competition among voice and subscription video providers, and may lead to increased deployment of broadband networks.[83]"

The italicized text might be read to suggest that if the FCC has classified some service in some other way, then it may not be considered a specialized service. This narrower use of the terms is made explicit in the merger agreement between Comcast and NBCU, which defines specialized service as follows:

'"Specialized Service" means any service provided over the same last-mile facilities used to deliver Broadband Internet Access Service other than (i) Broadband Internet Access Services, (ii) services regulated either as telecommunications services under Title II of the Communications Act or as MVPD services under Title VI of the Communications Act, or (iii) Comcast's existing VoIP telephony service[84].'

[83] Federal Communications Commission, Notice of Proposed Rulemaking, In the Matter of Preserving the Open Internet Broadband Industry Practices, GN Docket No. 09-191, WC Docket No. 07-52, FCC 09-93, (October 2009) 148 [italics added, footnote omitted]

[84] Federal Communications Commission, In the Matter of Applications of Comcast Corporation, General Electric Company and NBC Universal, Inc. For Consent to Assign Licenses and Transfer Control of Licensees, MB Docket No. 10-56, FCC 11-4, Appendix A, I (Definitions), pg. 121

This text makes explicit that in the context of the Comcast-NBCU Order, specialized service does not include Title VI MVPD service. Yet the R&O states that that IP video is explicitly included in the definition (but also, as noted above, may exclude services that are otherwise classified). These varied definitions have slowed the working group's progress, and may require future clarification by the FCC.

For the purpose of this working group, which functions in the context of the R&O, we have attempted to work with the inclusive definition of specialized service. The term as we use it is thus only meaningful within the context of the R&O. Used in this way, "Specialized services" are not a new category of items for regulation. Rather, they set a limit on which IP-based services are subject to the Open Internet rules. In this usage, some specialized services, such as VoIP and video, may already be subject to regulation under other laws and orders – the Open Internet R&O does not affect these other regulations. Rather, the labeling of a service as "specialized" would mean that that service is not subject to further regulation under the R&O.

We proceed with this definition, mindful of the fact that all such use of the term should properly be prefaced with OI, as in "OI specialized services".

Criteria for distinction

Based on the reading of the R&O, and subsequent discussions with FCC staff counseling the OIAC, there are two criteria in the R&O that would move a managed service far enough away from the open Internet that the R&O would not apply.

1) The service is not used to reach large parts of the Internet.
2) The service is not a generic platform but a specific "application level" service.

Using a number of case studies, we tried to tease out other aspects of a service that would set it apart from the services covered by the rules of the R&O. We identified one other criterion that we bring to the attention of the FCC.

1) Capacity isolation. The criterion of "capacity isolation" came up in a number of working group case studies, including the IPTV case study, the third-party platform case, and VoIP[85]. The argument is that a specialized service should not take away a customer's capacity to access the Internet. Since statistical multiplexing among services is standard practice among network operators, the isolation will not be absolute in most cases. However, if a specialized service substantially degrades the BIAS service, or inhibits the growth in BIAS capacity over time, by drawing capacity away from the capacity used by the BIAS, this would warrant consideration by the FCC to further understand the implications for the consumer and the possible competitive services running on the BIAS service.

Distinctions between BIAS and specialized services

[85] Voice over IP, or VoIP, is not a case study elaborated in this report, but was discussed by the working group, and shares the isolation attributes of IPTV.

The discussions concerning the differences between specialized services and a BIAS service tend to focus on the fact that specialized services, since they are not bound the requirements of the R&O, can offer different sorts of services, in particular enhanced service qualities. However, there will be other dimensions along which the services may differ; providers of BIAS who have usage tiers or usage caps need not impose those caps on specialized services, and specialized services may be priced and packaged in different ways.

High-level principles

We identified three high-level principles that the FCC should consider if and as it further deliberates about specialized services:

- Open Internet regulation should not create a perverse incentive for operators to move away from a converged IP infrastructure. Using IP should not imply a regulatory burden related to any regulation of the Internet.
- A service should not be able to escape regulatory burden, or acquire a burden, by moving to IP. A service may change or evolve as it migrates to IP, and the regulatory implications of such a change should be evaluated based on its characteristics.
- Proposals for regulation should be tested by applying them to the range of technologies being used for broadband. To the extent possible, regulation should be technology-neutral. (There are painful edge-conditions to this principle, which we acknowledge.)

These seem like simple statements, but in fact they may have very powerful consequences. They are an attempt to bound the scope of regulation without the need to debate the definition of any terms such as specialized services.

Monitoring the Internet

In recognizing specialized services as a category that is not subject to the Open Internet rules, the FCC also expressed the importance of ensuring that specialized services do not deter or limit investment in Internet services. The FCC expressed concern that "broadband providers may constrict or fail to continue expanding network capacity allocated to broadband Internet access service to provide more capacity for specialized services."[86] The FCC has declared their intention to monitor this situation. This committee is asked to advice them as to how to undertake this task.

Two approaches may address these concerns, although neither approach is wholly satisfactory and both approaches carry the risk of unintended consequences. On the one hand, the FCC may choose to define how much Internet service is "enough", and compare actual offerings to this standard. By setting a minimum standard for how much capacity for Internet service is available, the FCC could potentially make sure that sufficient capacity exists for providers of high-level service to innovate. It is important to note, however, that this minimum standard would likely have to change over time as consumers' usage habits and expectations shift. Alternatively, the FCC could compare what innovators can do using a specialized service as compared with the public Internet. Such a comparison would help the FCC to determine whether ISPs are exploiting

[86] R&O at 61.

a significant set of innovative opportunities via specialized services that are not available to others who would like to innovate over the open Internet. This second approach would reveal not only raw capacity concerns, but also quality of service concerns. As is illustrated in the third-party platform case study, the issue of comparing what can be done over the Internet and as a specialized service is not a simple matter of capacity, but depends on several parameters of the service.

The FCC currently performs a range of measurements on the Internet, tracking metrics such as achieved throughput, latency, and so on. In our discussions of specialized services, we did not identify any additional technical metrics that might be usefully measured, in order to better understand the impact of specialized services on the BIAS service. Instead, we focused on the higher-level question of what to make of these measurements—what sort of results would lead to the conclusion that the Internet was "good enough".

Exploration of this question is our tentative task for the next study period, but we have identified a possible approach to the issue. We believe that a promising approach is to start by looking at the quality of the user experience, not the technical parameters. The National Academies, in a 2002 report titled "Broadband: Bringing home the bits"[87], chose not to define broadband in numerical terms, because the committee knew that the target number would change over time. Instead, they defined it in terms of the needs of the applications of the time. They offered two definitions: a baseline definition and a forward-looking definition.

- Broadband Definition 1. Local access link performance should not be the limiting factor in a user's capability for running today's applications.
- Broadband Definition 2. Broadband services should provide sufficient performance and wide enough penetration of services reaching that performance level to encourage the development of new applications.

Neither definition is quantified, and neither, as stated, could directly be used as the basis of regulatory specification. However, the view of the committee was that these definitions could be translated into numbers that would be applicable at a given time. Based on our initial discussions, we believe that there have been a number of studies that relate the various technical parameters describing broadband performance to the operation of specific applications. We plan to explore this (and potentially other) approach to answering the question of when an Internet service is "good enough".

[87] Computer Science and Telecommunications Board, *Broadband: Bringing Home the Bits*, National Academy Press, 2002.

Appendix 1: Case study of IPTV

The Specialized Services working group is examining a range of issues surrounding "specialized services" in the context of the Open Internet Order, and how they relate to broader Internet access service and innovation. This appendix looks at the role of video (including IP based video) services, in today's marketplace and the potential effects on broadband Internet access service (BIAS). The paper provides a high-level overview of certain access network architectures, describes how services can be delivered over those architectures, and then discusses possible implications for BIAS.

High level overview of broadband access network architectures

Broadband Internet networks typically have a common general structure: the network operator's backbone connects to the networks of other operators and to its regional metro network, which in turn connects to local access facilities all of which contain fiber, optical components, routers, servers, switches and the like. The focus of this paper is on the access network, which is the portion of the network closest to the customer, and most relevant to the provision of specialized services over a shared facility that is used to deliver BIAS. Access networks typically comprise a mix of fiber and either coaxial cable (cable systems) or copper facilities (telco) to the home, and more recently, some network providers are using fiber facilities all the way to the home. Modern cable systems typically use a Hybrid Fiber Coax (HFC) access network, while telecommunications service providers typically use either a Digital Subscriber Line (DSL) or Passive Optical Networking (PON) based technology.

In a typical implementation of an HFC system, a cable operator will extend fiber from a Cable Modem Termination System (CMTS) to an Optical Node in a local neighborhood, which can serve anywhere from a few to several hundred homes. From each Optical Node, coaxial cable is then used to deliver service to the home. Services are delivered over Radio Frequency (RF) over coax typically using frequency bands from 52 MHz to 1000 MHz for downstream signals and 5 MHz to 42 MHz for upstream signals. The IP bandwidth is delivered by bonding together multiple 6 MHz RF channels, the same channels that traditionally were used to deliver a single analog video channel (explained later). With the recent DOCSIS 3.0 specification, cable operators typically bond 8 channels downstream to support a downstream channel of approximately 300 Mbps (although some cable operators are starting to bond 12 and 16 channels downstream), which is shared among a number of subscribers attached to a given node. Depending upon the details of the HFC infrastructure, the total number of subscribers connected to an Optical Node, and the number of subscribers online at a given point in time, this architecture can deliver a wide range of BIAS speeds along with specialized services.

Telecommunications service providers have typically used DSL and more recently PON systems to deliver service to the home. Similar to cable operators, over time DSL providers have extended fiber optics closer to homes, using some combination of Fiber to the Node (FTTN) and Fiber to the Home (FTTH). The emergence of next generation DSL technologies, such as Very

High Rate DSL (VDSL), ADSL2+ and techniques such as pair bonding and vectoring have enabled service providers delivery speeds much higher compared with legacy DSL technologies. In the case of a DSL implementation, the broadband connection in the access network is dedicated to an end user from the node to a user's home, rather than being shared as in typical cable HFC systems. For FTTH implementations, most service providers are using a technology referred to as Passive Optical Networks (PON). PON systems generally take one strand of fiber to a fiber splitter location, and then replicate the optical signal onto multiple separate fiber strands connected to subscriber homes. A PON system consists of an Optical Line Terminal (OLT) placed in a serving central office and an Optical Network Terminal (ONT), or electronics, at the subscriber premises. As with VDSL services, this technology can deliver speeds far in excess of traditional DSL.

Service delivery methods

Services delivered over these architectures typically include video, voice, and BIAS services. Broadband providers offering video services are classified as Multichannel Video Programming Distributors (MVPDs). Different MVPDs deliver video service in a variety of ways. Most cable systems today, and in some instances PON based video services, provide live linear programming ("traditional TV") using specific frequency bands dedicated to specific channels. All channels are simultaneously delivered or "broadcast" to the subscriber's premises, and tuners in the set top box act as filters to permit display of the desired programming network. For Video on Demand (VOD) services, MVPDs typically dedicate certain channels for delivery of requested content. In some cases, cable operators are offering linear programming networks and VOD delivered using IP or another packet-based transmission system, however, the vast majority of live linear video programming continues to be delivered using specific frequency bands dedicated to specific programming networks.

Modern cable systems use a digital representation of video, either compressed Motion Picture Expert Group (MPEG)-2, or more recently MPEG-4, video modulated onto Quadrature Amplitude Modulated (QAM) RF signals. In a typical implementation, a cable operator will organize the bandwidth used for digital video into the same 6 MHz channels of frequency as it would in a traditional analog cable system and, using 256 QAM, deliver approximately 38 Mbps per 6 MHz channel. In a typical MPEG-2 configuration, a Standard Definition (SD) channel can be encoded in a range from 2-6 Mbps and High Definition Content ranging from 15-19 Mbps. MPEG-4 halves these ratios to around 2-3 Mbps for SD and 6-7 Mbps for an HD channel. Thus a single 6 MHz channel slot with 256 QAM at approximately 38 Mbps could deliver up to 2 HD channels or 10 SD channels with MPEG-2, or perhaps twice that capacity with MPEG-4. The High Efficiency Video Encoding (HVEC) currently under development by the ISO/IEC Moving Picture Expert Group (MPEG) and the ITU-T Video Encoding Expert Group is intended to be the successor standard to MPEG-4 and is projected to reduce the bandwidth requirement by 50% for the same quality picture. It can also support resolutions up to 8192x4320.

The BIAS services offered over these cable systems will typically use separate and distinct channels and frequencies from the linear video services, creating a separation between the services sharing the infrastructure and dedicating fixed amounts of bandwidth to each service. As noted, in some markets cable operators have begun offering traditional cable video services, both linear channels as well as VOD, in IP format. These IP cable services are delivered over the IP bandwidth a cable operator creates by bonding multiple 6 MHz channels, but these IP cable services typically use a separate service flow to customers' homes – with bandwidth above and beyond the bandwidth allocated for the customer's BIAS service – that is allocated specifically for the IP cable service

Another means of service delivery is a pure IP based infrastructure where all services are carried using IP on the same physical network. In this case, all video will be carried as IPTV. Any broadband IP network, regardless of the access network infrastructure, can be used for IPTV. The continuous improvements in data transfer speeds, brought about by advancements in both Digital Subscriber Line (DSL) and cable DOCSIS technology, combined with the improvements in compression ratios (e.g. the greater bandwidth efficiency offered by MPEG-4 over MPEG-2), and the emergence of switched digital video have enabled more video streams at higher quality to be delivered over broadband than previously possible.

The broadest use of IPTV has been by telecommunications operators to enable video delivery over their existing copper loop infrastructures. In contrast to broadcast video distribution typically used by cable companies, IPTV services utilize a switched, two-way, client server based architecture. Thus when a user "tunes in" to a "channel" delivered by an IPTV system, they are actually sending a request to initiate a stream of IP packets containing the requested video, and the servers stream only the requested content.

Capacity isolation

As the previous discussion suggests, one factor that distinguishes different methods of delivering services is how the overall capacity of the physical access path is allocated to the different services. On cable systems, the capacity used for traditional video (encoded over QAMs) is separate from the capacity for BIAS. When the video service migrates to IPTV, the capacity that is allocated to the IPTV service may be isolated from the BIAS capacity to different degrees. In general, IP bandwidth to the home is dynamically allocated, meaning that varying amounts of bandwidth will be allocated to different services, depending upon the exact network usage of the household at a given moment in time.

Different technologies may accomplish capacity isolation among services in different ways. Cable systems using DOCSIS may open a separate service flow for the MVPD IPTV and allocate capacity to that flow sufficient for the video. In this way, the possibility that the IPTV and the BIAS may affect each other is minimized. On some other systems the allocation of capacity between MVPD IPTV and BIAS may not be as rigid. Based on information from the members of the subgroup familiar with current practices, most schemes for delivery of MVPD IP video attempt to isolate the capacity used for MVPD and BIAS to a high degree. However, public documentation is usually not specific as to practices.

The previous discussion has focused on the access path into the residence, but issues of traffic isolation can also arise in other parts of the network. Depending where the content servers are, the IP traffic between the servers and the access network might be totally segregated from the